Long Term Growth Rates: Estimation and Uncertainty

Jim Stock

Harvard University

Social Security Technical Panel on Assumptions and Methods

March 29, 2019

Outline

1. Growth identities
2. Estimating uncertainty of long-run averages
3. Three methods for estimating trends
4. Trends and confidence intervals for 75-year mean growth (selected series)
5. Improvements & extensions
6. Digression on long-term real rate
7. Summary

Growth identities

GDP per capita

$$\Delta \ln\left(GDP_t / Pop_t\right) = \Delta \ln\left(GDP_t / Hr_t\right) + \Delta \ln\left(Hr_t / Emp_t\right) + \Delta \ln\left(Emp_t / LF_t\right) + \Delta \ln\left(LF_t / Pop_t\right)$$

$$y_t = \mu_t + h_t + z_t + l_t$$

Earnings per capita

$$\Delta \ln\left(Earn_t / Pop_t\right) = \Delta \ln\left(Earn_t / NGDP_t\right) + \Delta \ln\left(NGDP_t / GDP_t\right) + \Delta \ln\left(GDP_t / Pop_t\right)$$

$$e_t = s_t + \pi_t + y_t$$

Note: The OASDI Report (2018) uses these two identities but further breaks out earnings into (earnings/comp)×(comp/NGDP)

Summary - key variables for characterizing long-term growth rates:

$$x_t = \begin{pmatrix} \mu_t \\ l_t \\ s_t \\ h_t \\ z_t \\ \pi_t \end{pmatrix} = \begin{pmatrix} \Delta \ln\left(\textit{Economy-wide productivity}\right) \\ \Delta \ln LFPR_t \\ \Delta \ln LaborShare_t \\ \Delta \ln WeeklyHrs \\ \Delta \ln(1-u_t) \\ \text{GDP deflator inflation} \end{pmatrix}$$

Estimating uncertainty of long-run averages

We are interested in the long-run average,

$$\bar{x}_{T+1:T+h} = \frac{1}{h}\sum_{s=T+1}^{T+h} x_s$$

A. Simplest setup: ***x* is stationary** (no trends or drifts). Then by the Central Limit Theorem,

$$\bar{x}_{T+1:T+h} \sim N\left(\mu, h^{-1}\Omega\right)$$

where Ω is the long-run variance of *x* :

$$\mathrm{var}\left(\bar{x}_{T+1:T+h}\right) = \mathrm{var}\left(\frac{1}{h}\sum_{s=T+1}^{T+h} x_s\right) \to h^{-1}\Omega = h^{-1}\sum_{u=-\infty}^{\infty}\Gamma_u, \text{ where } \Gamma_u = \mathrm{cov}(x_t, x_{t-u})$$

Monte Carlo approach to distribution of $\bar{x}_{T+1:T+h}$ with μ and Ω known:

$$\bar{x}_{T+1:T+h} = \mu + u, \quad u \sim N\left(0, h^{-1}\Omega\right)$$

Two problems: (1) μ isn't known, (2) Ω isn't known.

Estimating uncertainty of long-run averages

$$\bar{x}_{T+1:T+h} = \mu + u, \quad u \sim N\left(0, h^{-1}\Omega\right)$$

(1) Estimation of μ. The natural estimator of μ is $\bar{x}_{1:T} = \frac{1}{T}\sum_{s=1}^{T} x_s$. This adds estimation uncertainty:

$$\bar{x}_{T+1:T+h} = \mu + u = \bar{x}_{1:T} + \left(\mu - \bar{x}_{1:T}\right) + u = \bar{x}_{1:T} + v + u, \text{ where } v \sim N\left(0, T^{-1}\Omega\right) \text{ and } v, u \text{ are independent}$$

or

$$\bar{x}_{T+1:T+h} = \bar{x}_{1:T} + w, \text{ where } w \sim N\left[0, \left(\frac{1}{T} + \frac{1}{h}\right)\Omega\right]$$

(2) Estimation of Ω. This is the HAC problem of time series regression, so use a HAC estimator,

e.g., Newey-West. Putting this together, for a scalar *x*,

$$SE\left(\bar{x}_{T+1:T+h}\right) = \sqrt{\left(\frac{1}{T} + \frac{1}{h}\right)\hat{\Omega}}$$

This standard error captures two sources of uncertainty:

- Forecast uncertainty about what the future will hold: future random variation around μ
- Sampling uncertainty about the true mean μ

Estimating uncertainty of long-run averages

B. More complicated setup: small trend in x_t (local levels model):

$$x_t = \mu_t + \varepsilon_t$$

$$\mu_t = \mu_{t-1} + \eta_t$$

This produces the exponentially weighted moving average estimate $\mu_{T|T} = (1-\rho)^{-1}\sum_{i=0}^{T-1}\rho^i x_{t-i}$

Under the local levels model, the optimal prediction of $\bar{x}_{T+1:T+h}$ is $\mu_{T|T}$.

The prediction variance formula is more complicated but for small drift is approximated by the variance formula for the stationary case.

- I will use the local levels model to estimate μ, combined with the variance formula and HAC estimator of Ω for the stationary case

C. (Much) more complicated yet: allow for fractional integration and/or large AR components.

Műller & Watson (REStud, 2016): same idea, but a richer class of low-frequency models

Three methods for estimating trends

1. MA(40) (10 years, quarterly data)

$$x_T^{MA40} = \frac{1}{40}\sum_{i=0}^{39} x_{t-i}$$

2. Exponentially weighted moving average (EWMA):

$$x_T^{EWMA} = (1-\rho)^{-1}\sum_{i=0}^{T-1} \rho^i x_{t-i}$$

   - This is the smoother for the local levels model
   - ρ is chosen to have half-life of 10 years (not estimated)

3. Cyclically-adjusted biweight filter. Model: $x_t = \mu_t + \beta(L)u_t^{gap} + v_t$
   - 2-step kernel estimation of β(L):
   i. Deviate LFPR, u-gap from low-frequency trend (biweight kernel, BW = 40)
   ii. Regress deviated LFPR on deviated u-gap (*t*+2, *t*+1,…, *t*-8)
   - Trend estimate is biweight kernel applied to $x_t - \hat{\beta}(L)u_t^{gap}$

Note: I use (1) and (3) for plots, and (2) for the estimate of the 75-year average and for the center of the 67% confidence interval for the 75-year average. This allows for drift & trends.

Data, trends, and confidence intervals for 75-year mean growth

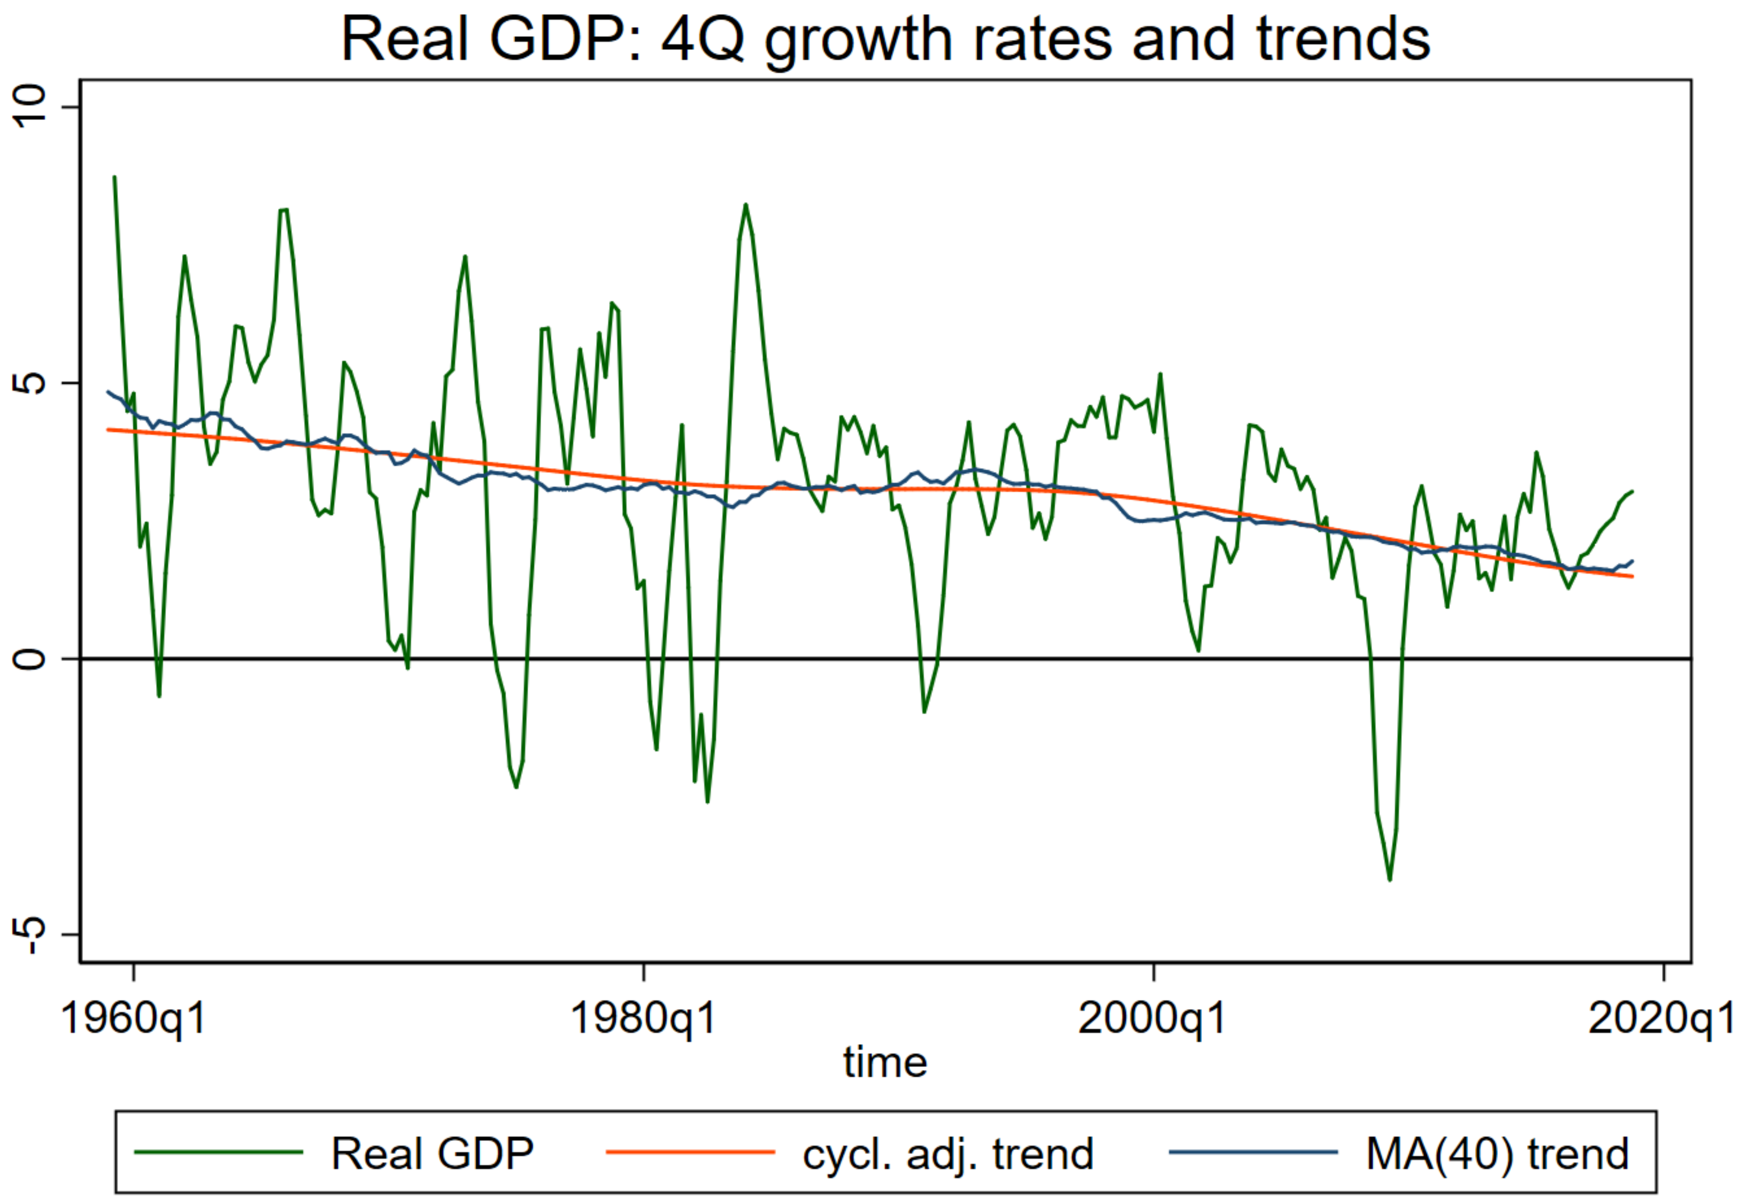

Sample mean, EWMA, and 75-year confidence intervals:

Statistic	%/yr
$\bar{x}_{1:T}$, 59-18	3.1
SE($\bar{x}_{1:T}$)	0.33
2019-2094	
EWMA, 2018:	2.4
67% CI	(2.0, 2.9)
90% CI	(1.7, 3.1)

CIs are computed using the Newey-West estimate of the long-run variance Ω, truncation parameter = 24, centered at the EWMA estimate of the long-run mean in 2018.

Data, trends, and confidence intervals for 75-year mean growth

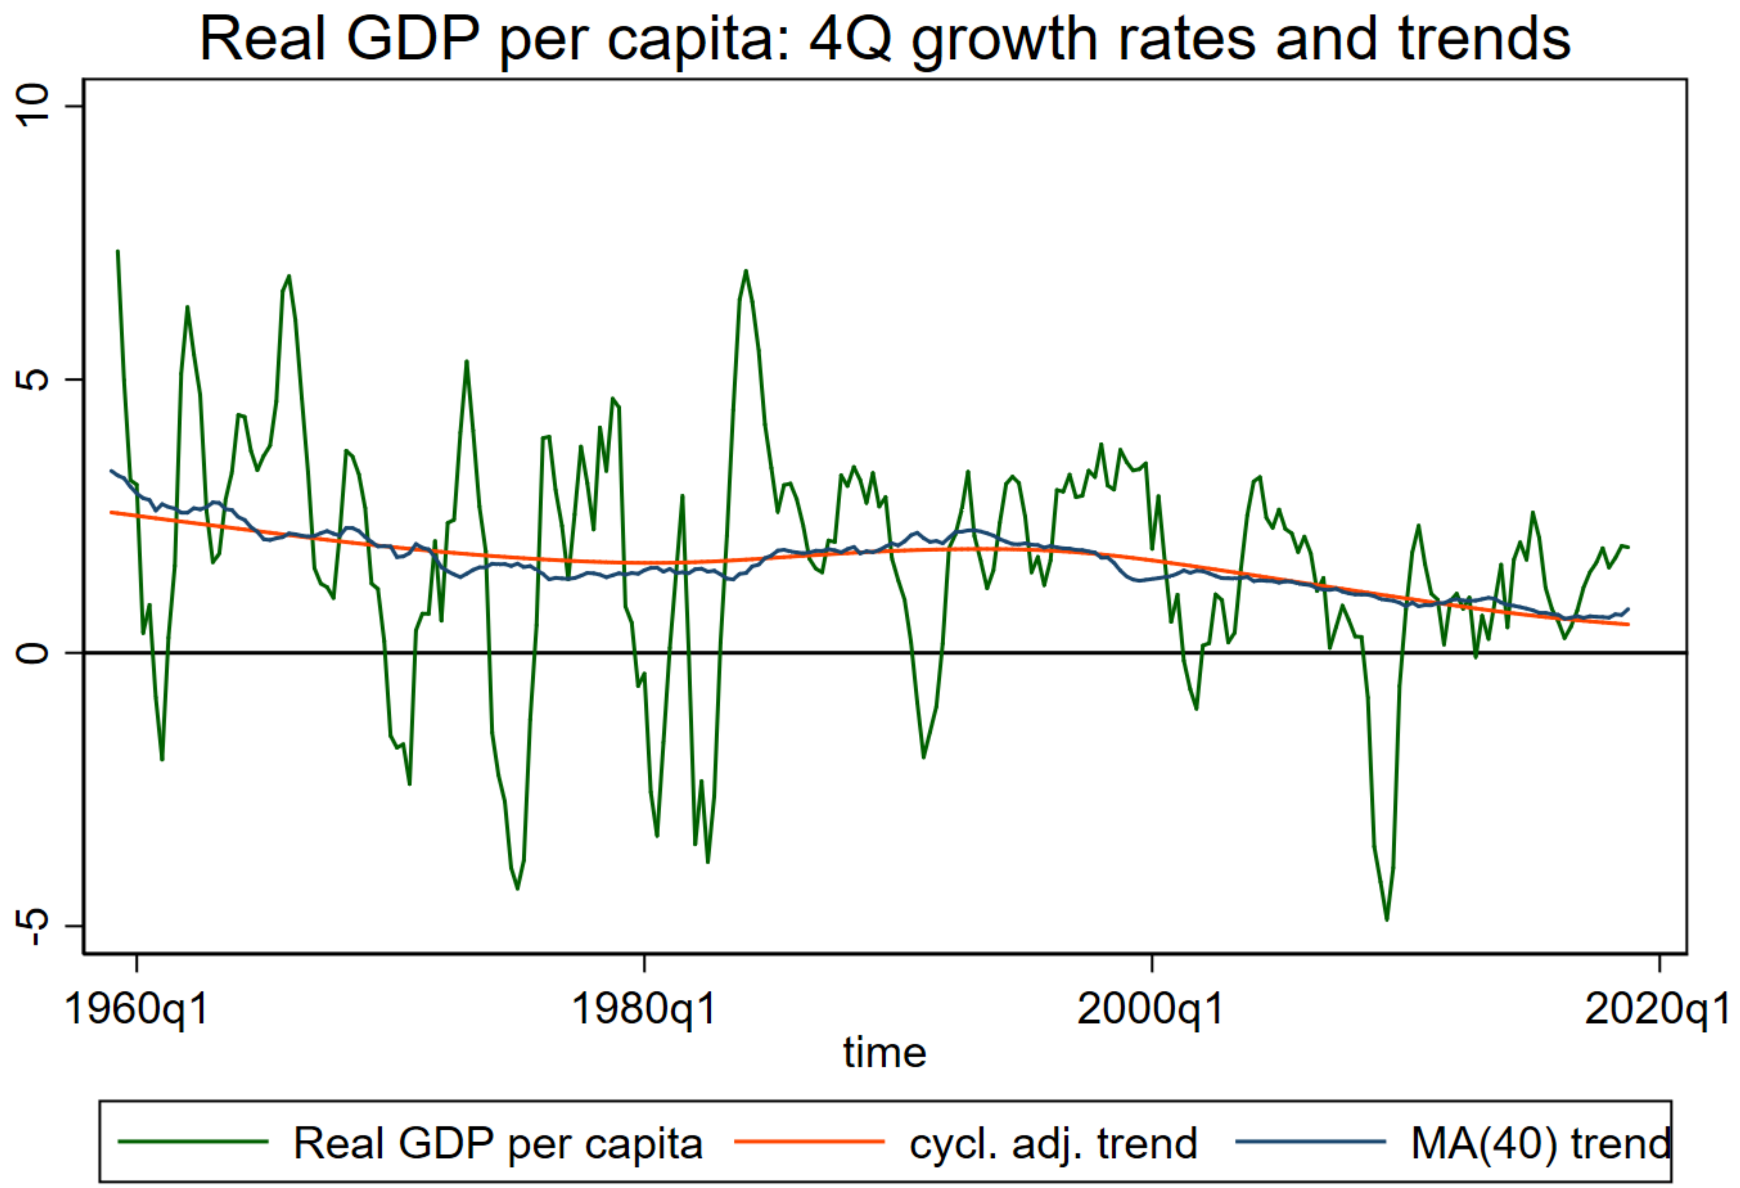

Sample mean, EWMA, and 75-year confidence intervals:

Statistic	%/yr
$\bar{x}_{1:T}$, 59-18	1.7
SE($\bar{x}_{1:T}$)	0.31
2019-2094	
EWMA, 2018	1.3
67% CI	(0.9, 1.7)
67% MW CI	(1.5, 2.4)
90% CI	(0.6, 2.0)

MW are Műller-Watson (2016) 67% confidence intervals using data from 1946-2014, 75-year projections

Data, trends, and confidence intervals for 75-year mean growth

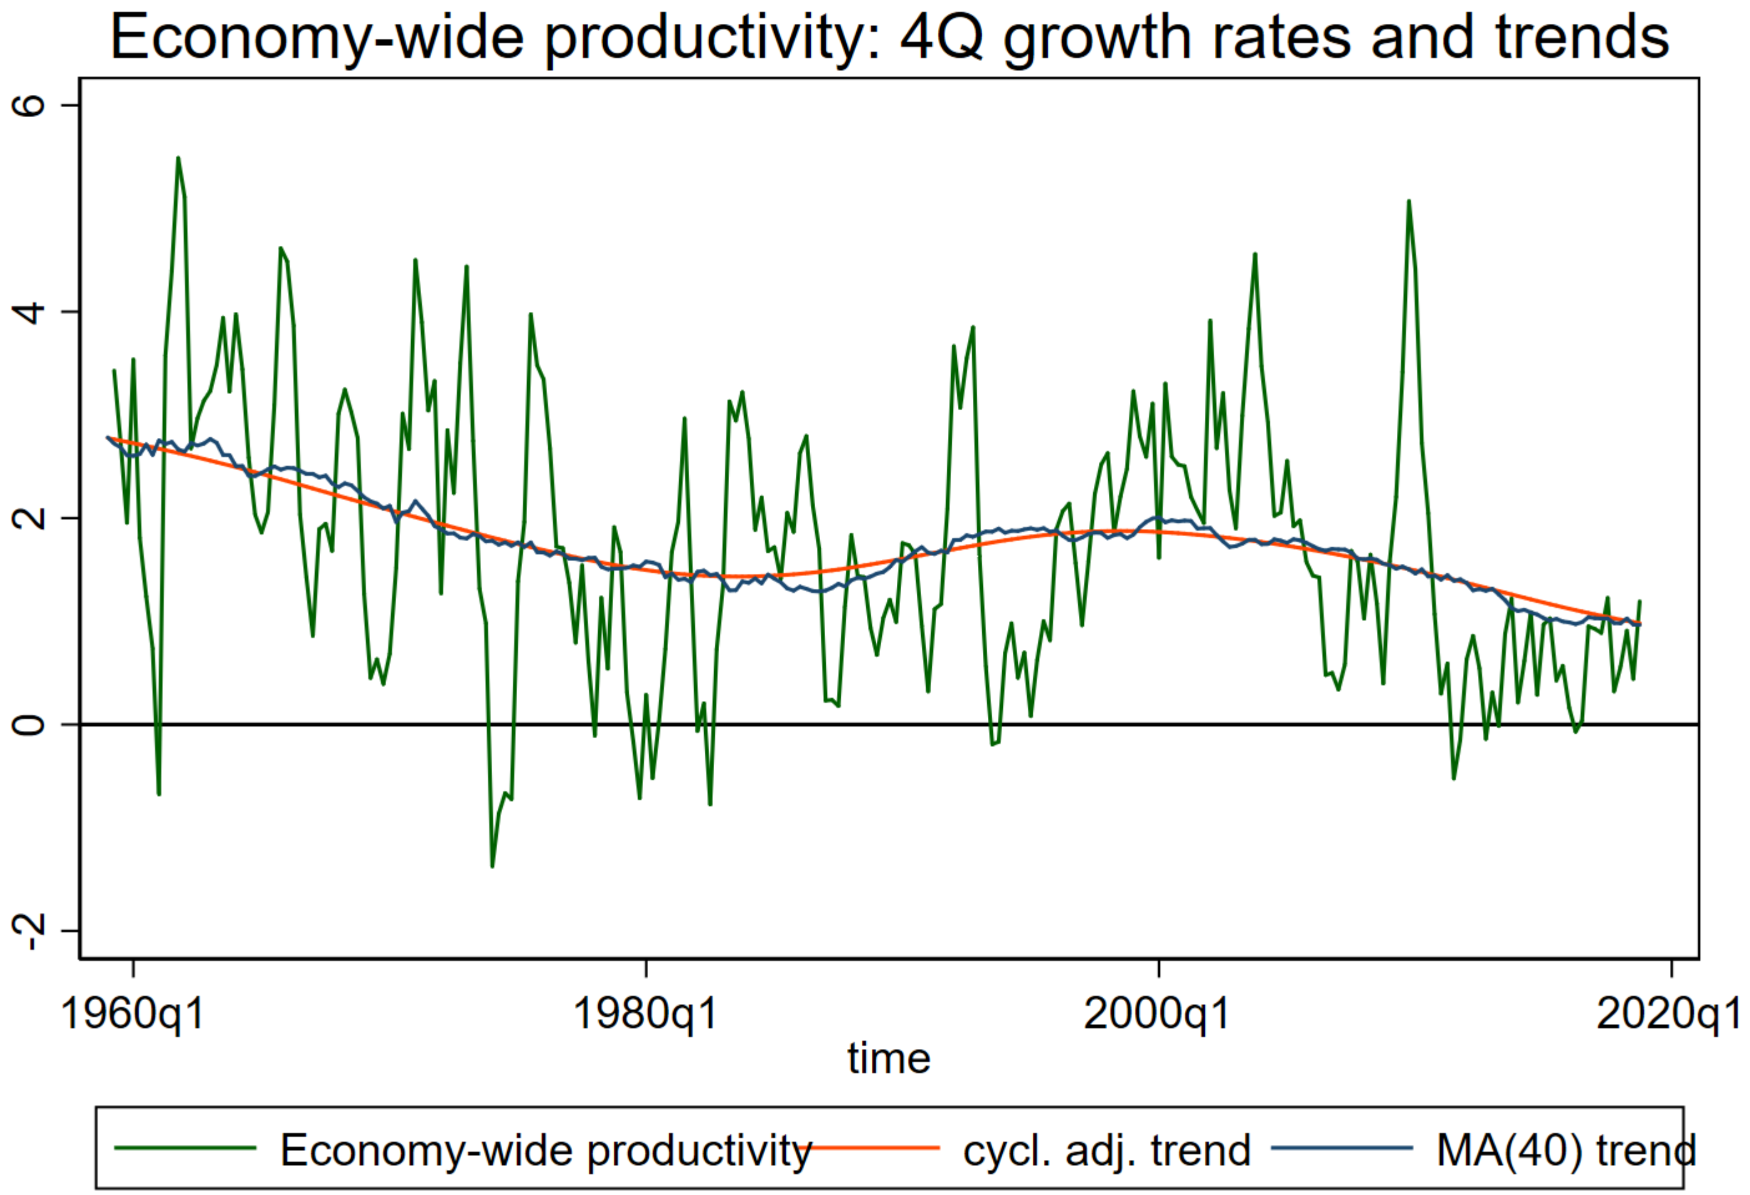

Sample mean, EWMA, and 75-year confidence intervals:

Statistic	%/yr
$\bar{x}_{1:T}$, 59-18	1.8
SE($\bar{x}_{1:T}$)	0.22
2019-2094	
EWMA, 2018:	1.4
67% CI	(1.1, 1.7)
90% CI	(0.9, 1.8)

MW are Műller-Watson (2016) 67% confidence intervals using data from 1946-2014, 75-year projections

Data, trends, and confidence intervals for 75-year mean growth

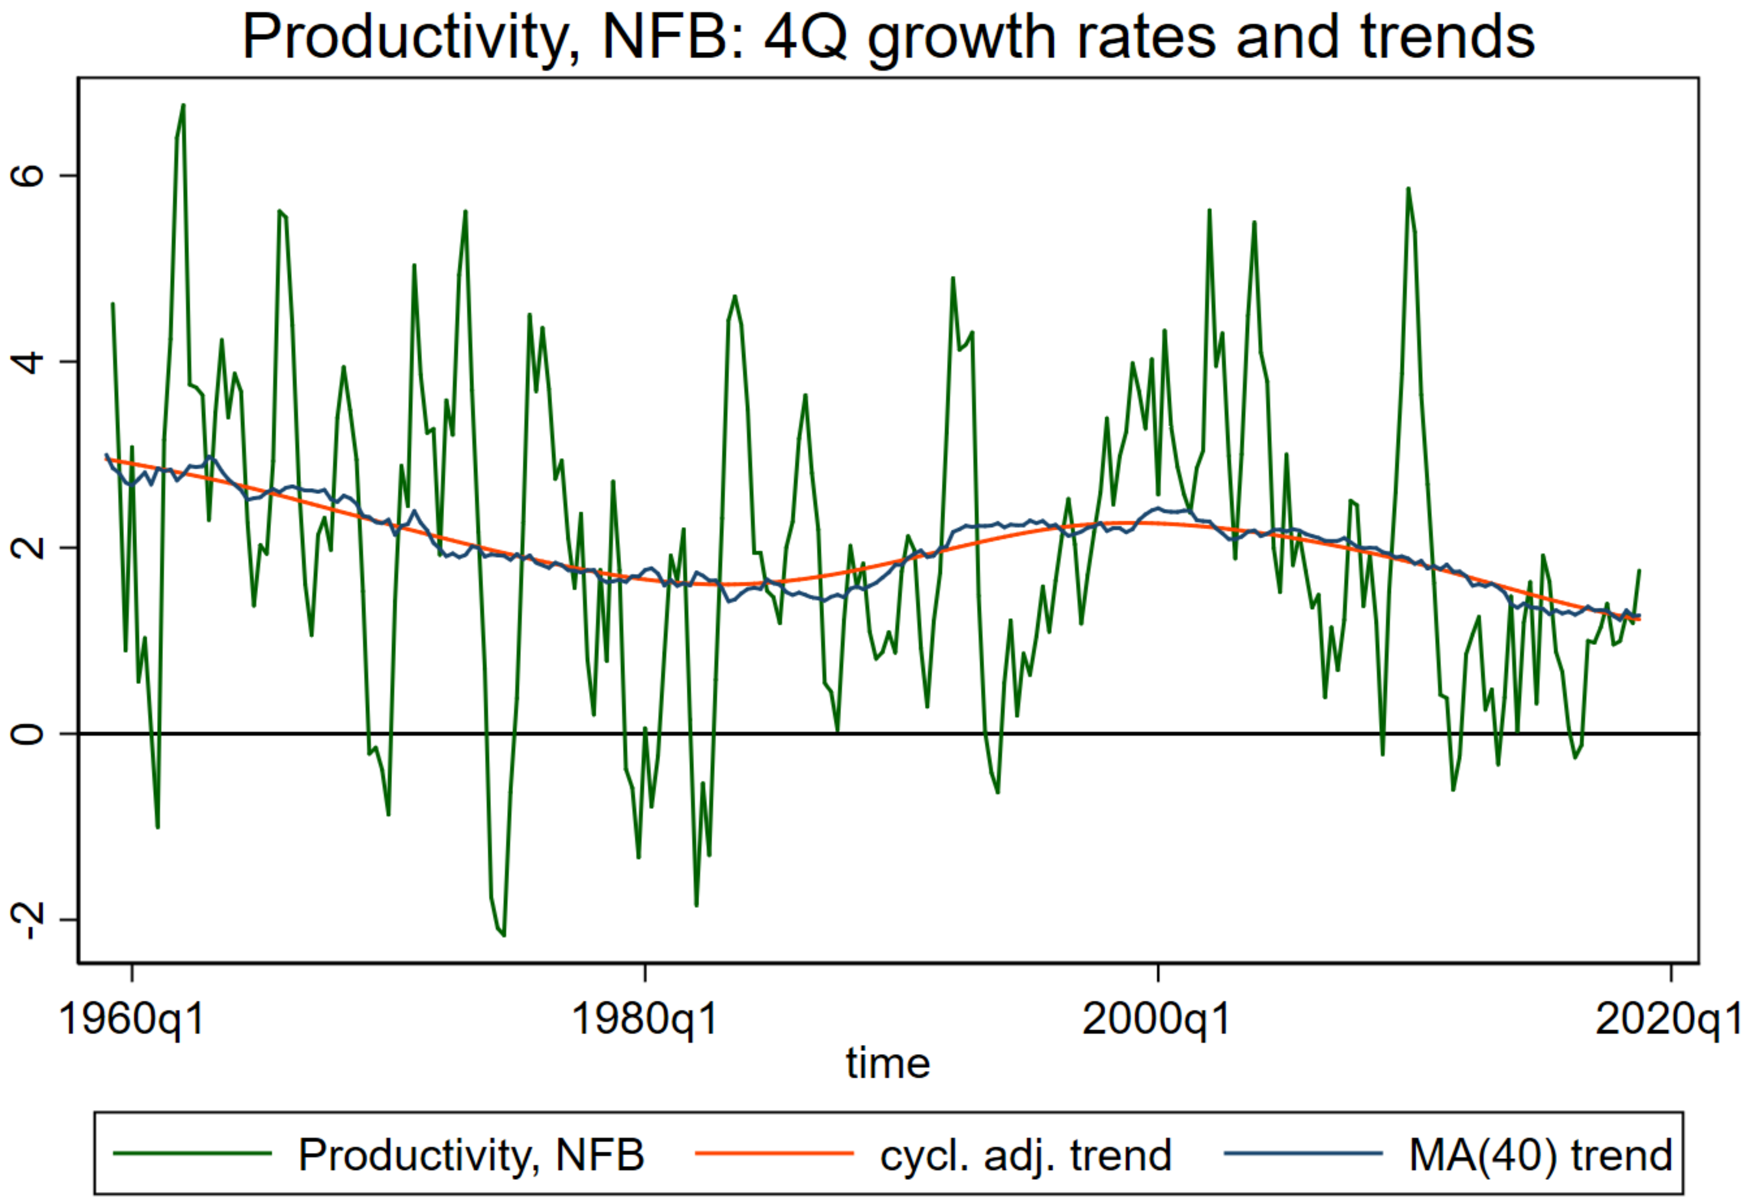

Sample mean, EWMA, and 75-year confidence intervals:

Statistic	%/yr
$\bar{x}_{1:T}$, 59-18	2.0
SE($\bar{x}_{1:T}$)	0.24
2019-2094	
EWMA, 2018:	1.7
67% CI	(1.4, 2.0)
67% MW CI	(1.3, 2.5)
90% CI	(1.1, 2.2)

MW are Műller-Watson (2016) 67% confidence intervals using data from 1946-2014, 75-year projections

Data, trends, and confidence intervals for 75-year mean growth

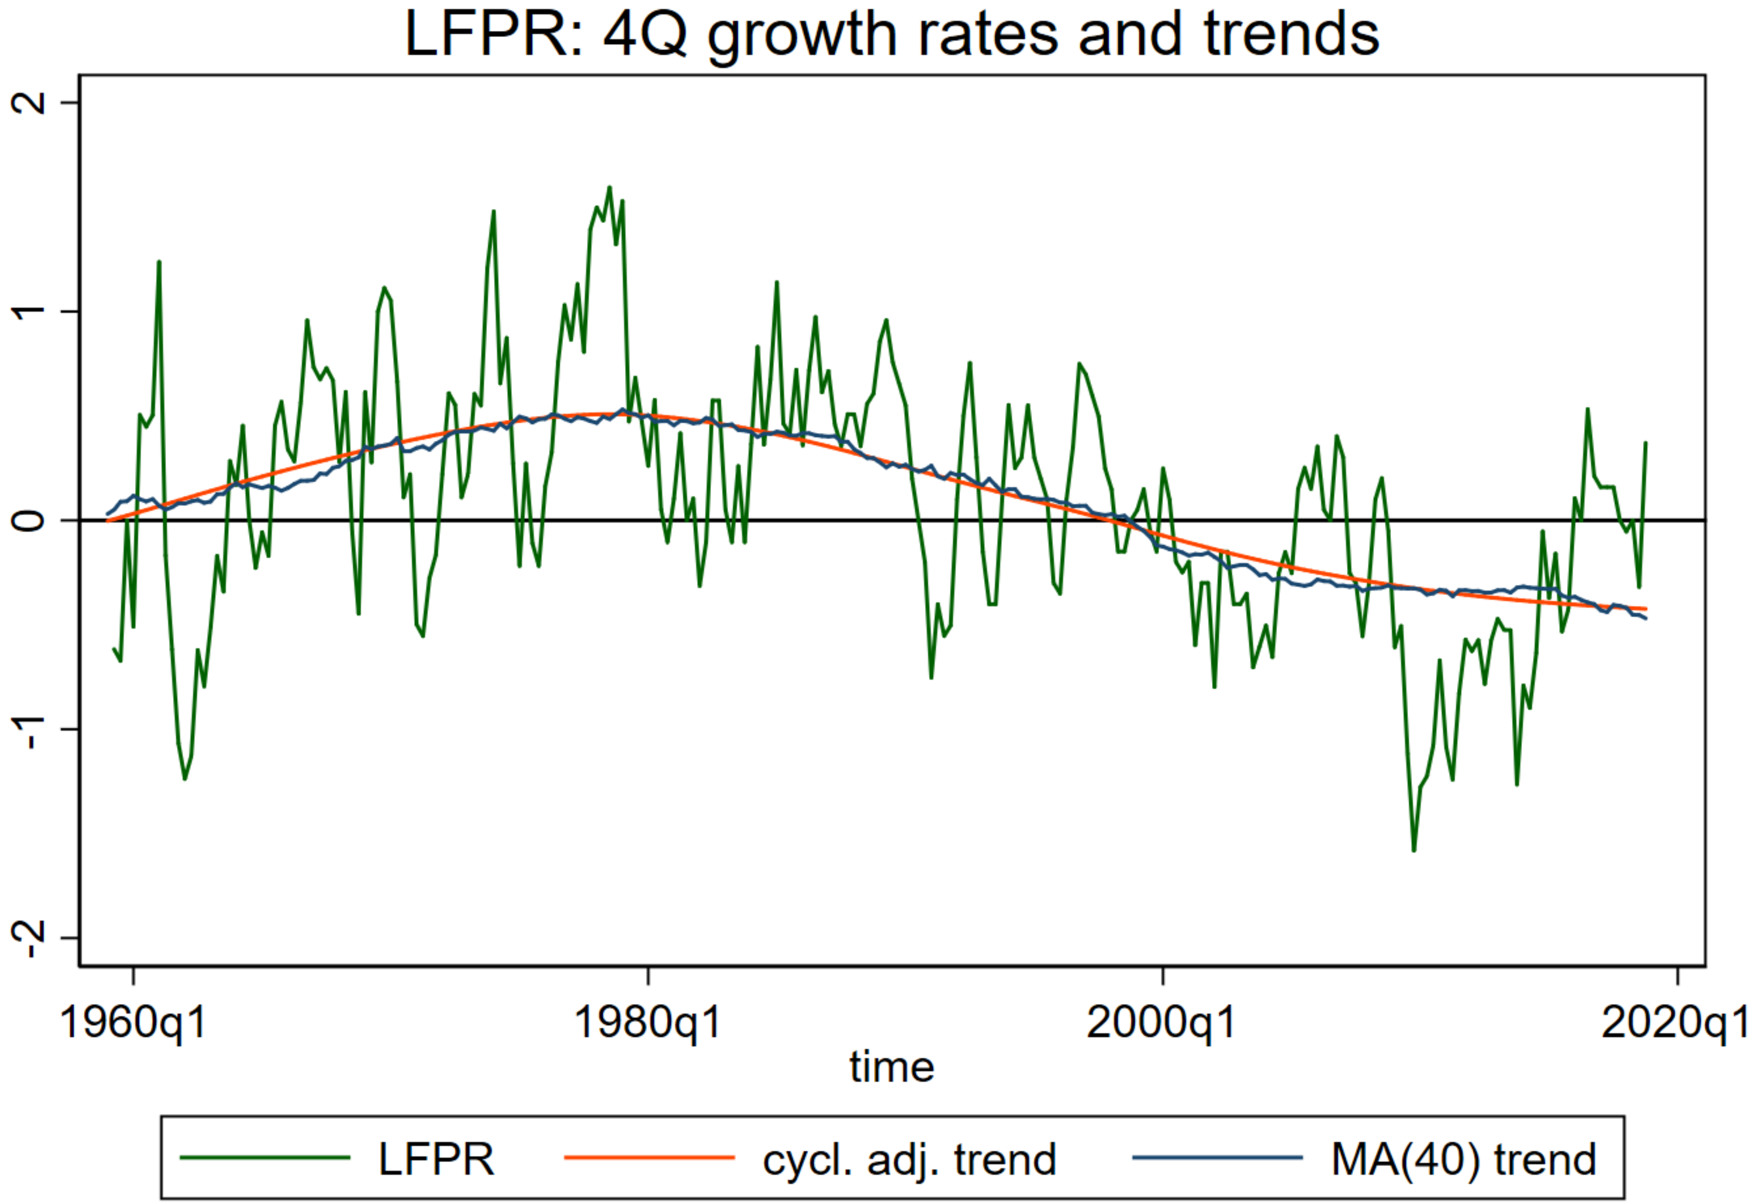

Sample mean, EWMA, and 75-year confidence intervals:

Statistic	%/yr
$\bar{x}_{1:T}$, 59-18	0.09
SE($\bar{x}_{1:T}$)	0.12
2019-2094	
EWMA, 2018:	-0.18
67% CI	(-0.34, -0.01)
90% CI	(-0.45, 0.10)

Data, trends, and confidence intervals for 75-year mean growth

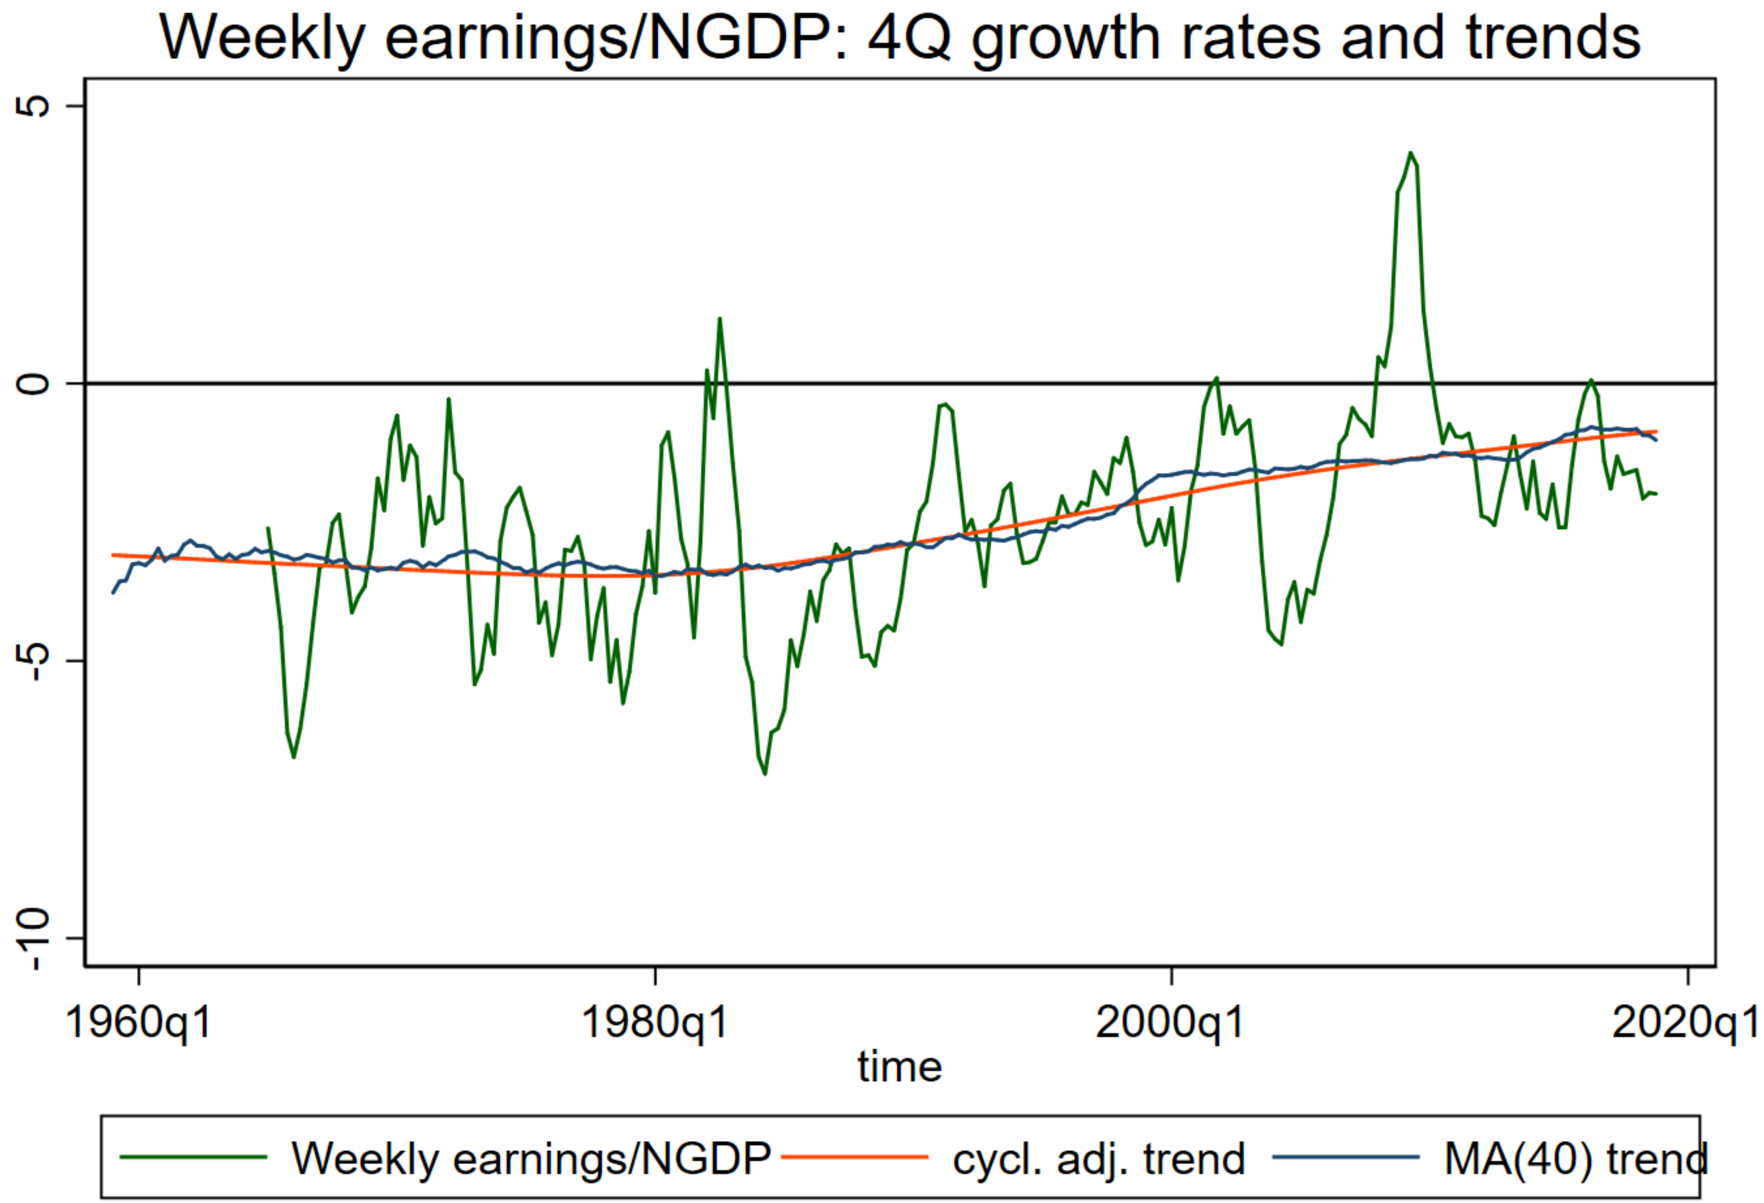

Sample mean, EWMA, and 75-year confidence intervals:

Statistic	%/yr
$\bar{x}_{1:T}$, 59-18	-2.5
SE($\bar{x}_{1:T}$)	0.36
2019-2094	
EWMA, 2018:	-1.7
67% CI	(-2.2, -1.3)
90% CI	(-2.5, -1.0)

Data, trends, and confidence intervals for 75-year mean growth

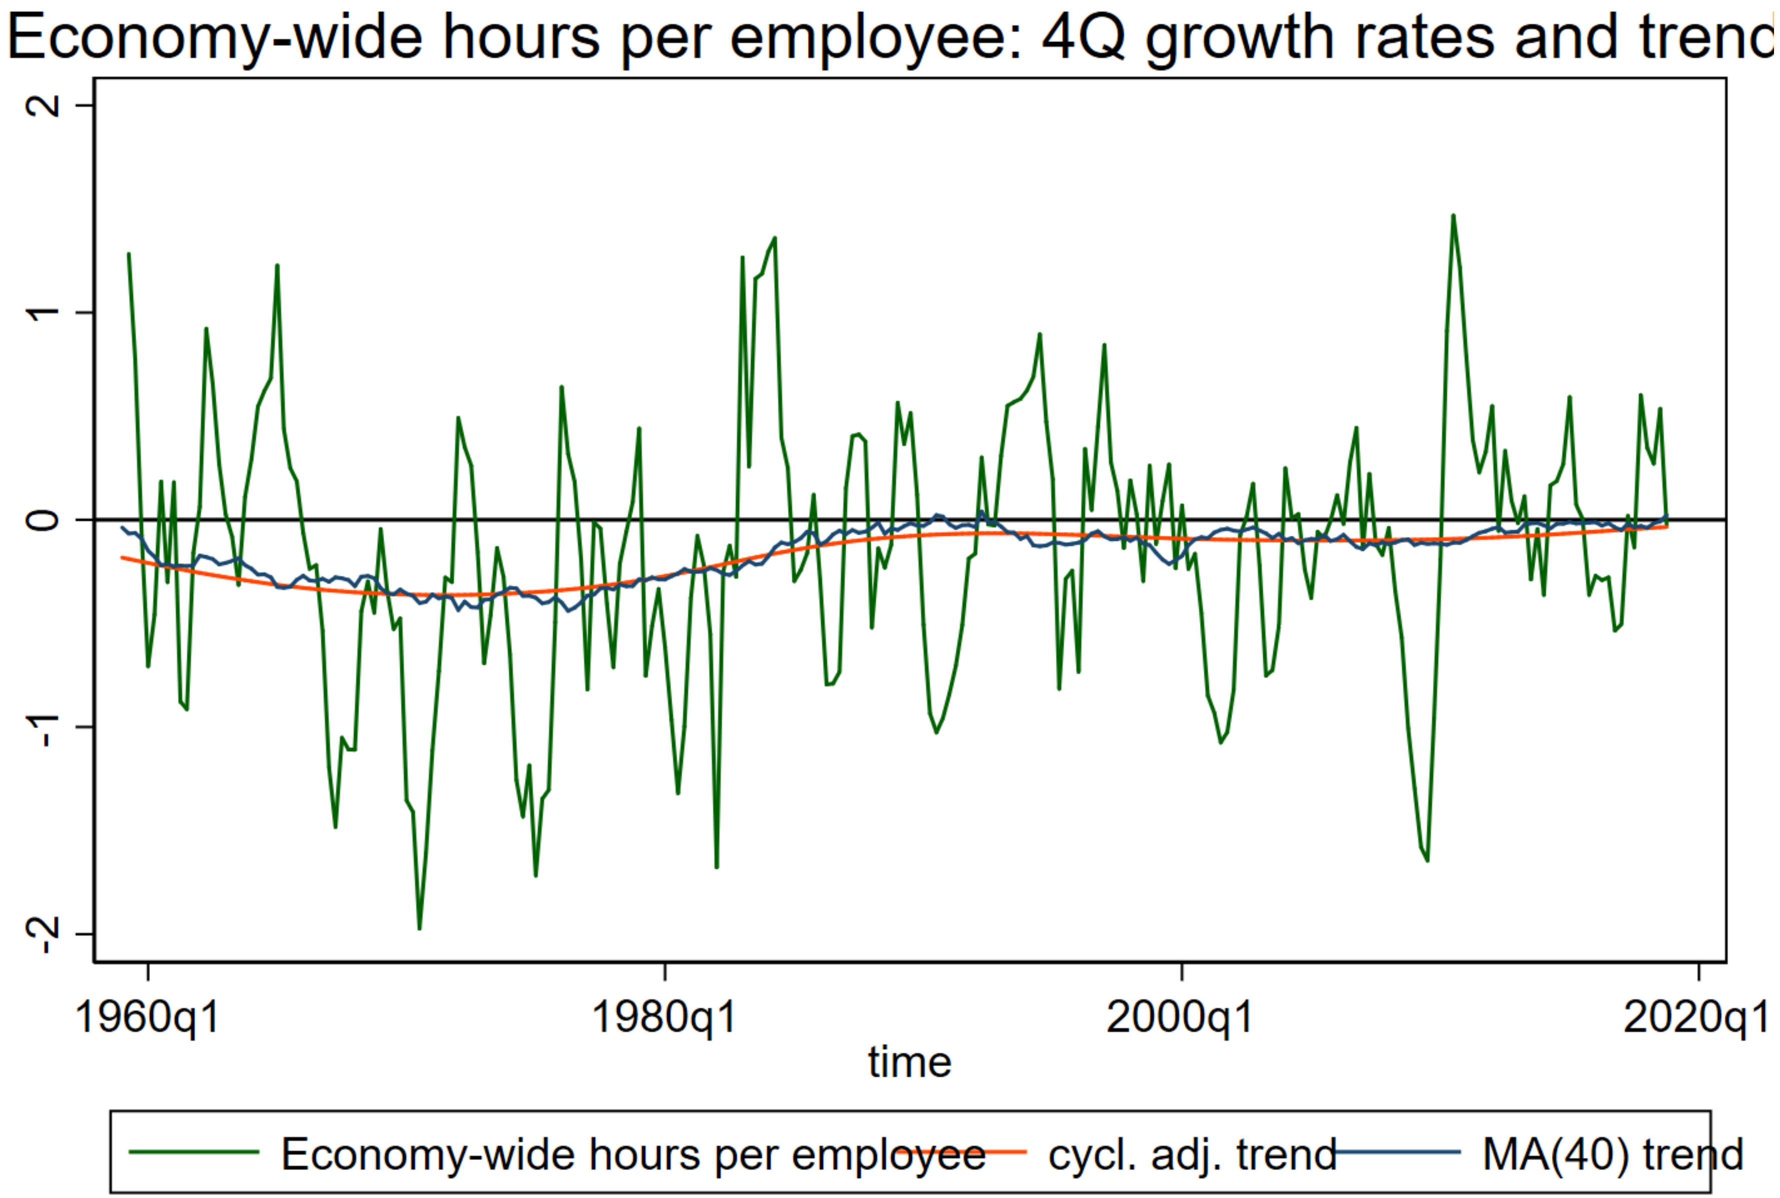

Sample mean, EWMA, and 75-year confidence intervals:

Statistic	%/yr
$\bar{x}_{1:T}$, 59-18	-0.14
SE($\bar{x}_{1:T}$)	0.08
2019-2094	
EWMA, 2018:	-0.05
67% CI	(-0.16, 0.06)
90% CI	(-0.23, 0.13)

Data, trends, and confidence intervals for 75-year mean growth

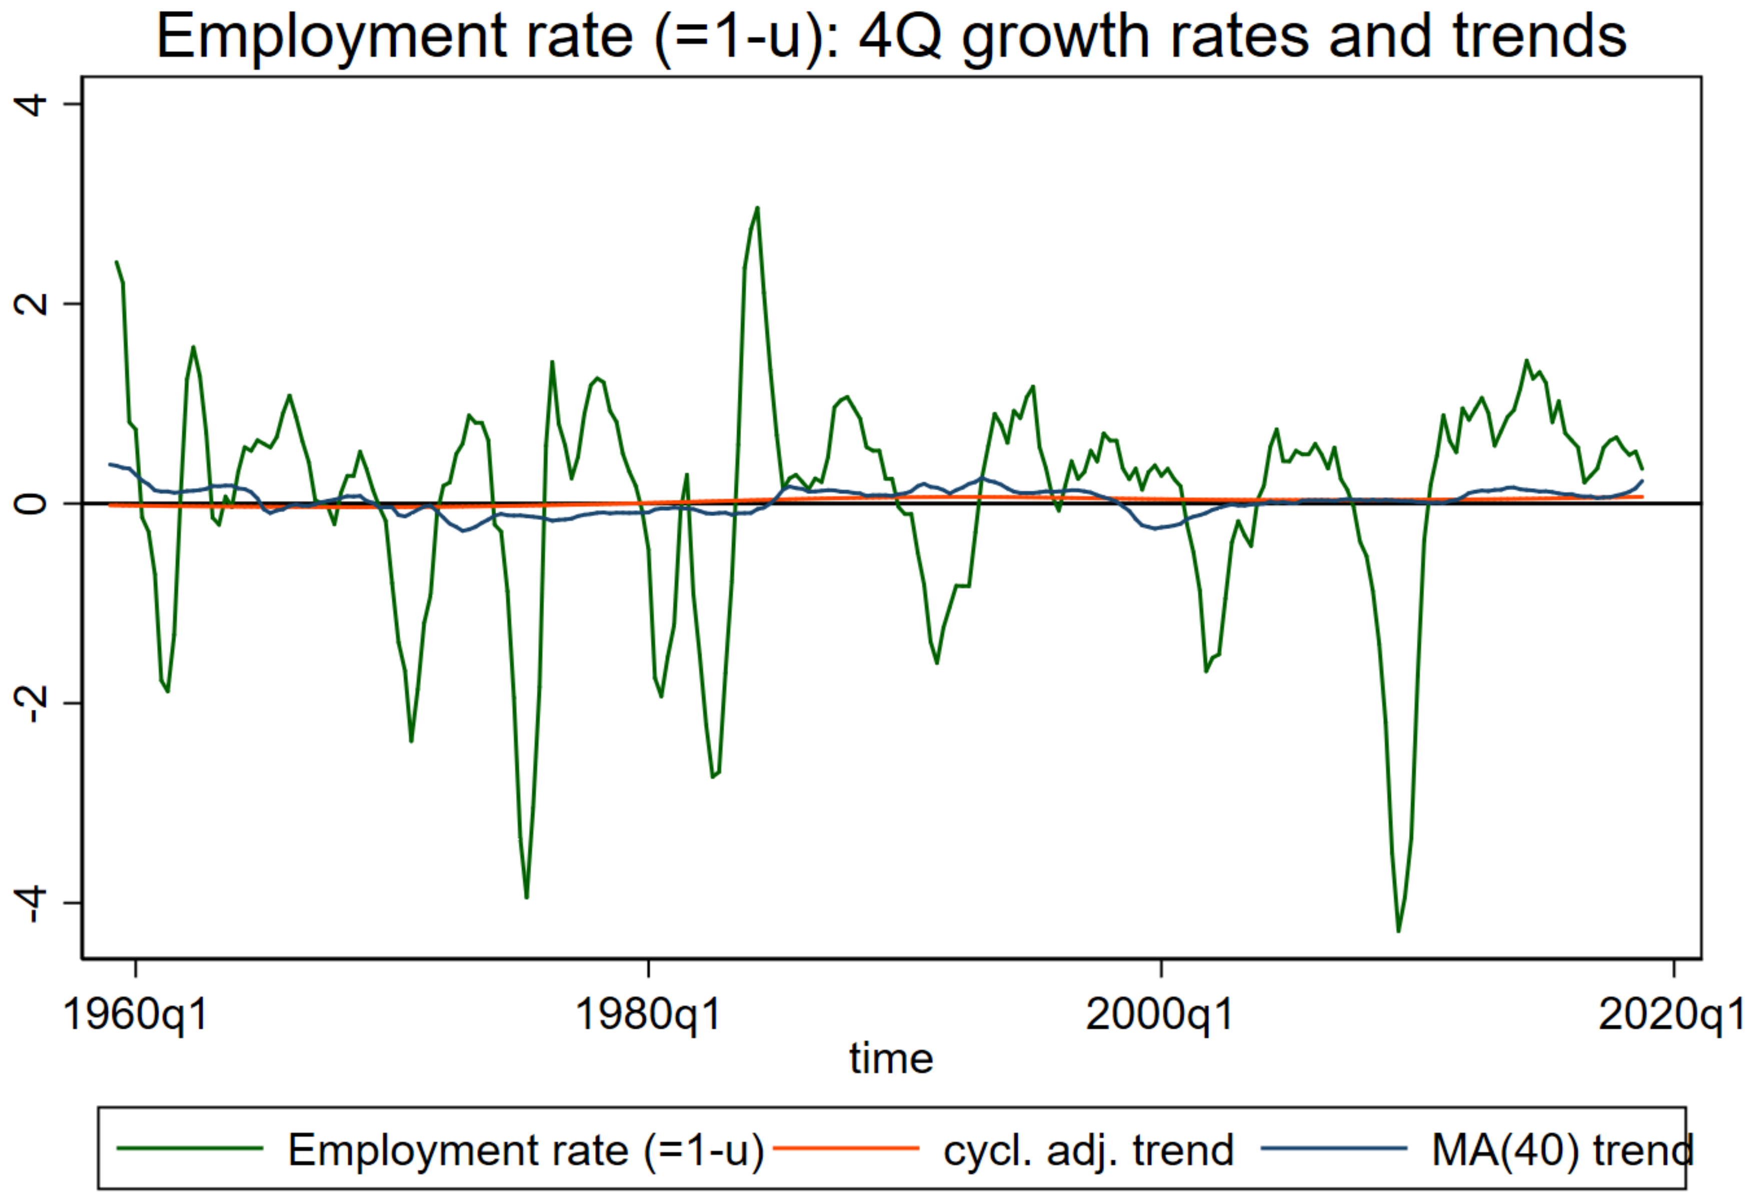

Sample mean, EWMA, and 75-year confidence intervals:

Statistic	%/yr
$\bar{x}_{1:T}$, 59-18	0.06
SE($\bar{x}_{1:T}$)	0.13
2019-2094	
EWMA, 2018:	0.17
67% CI	(-0.01, 0.35)
90% CI	(-0.12, 0.46)

Data, trends, and confidence intervals for 75-year mean growth

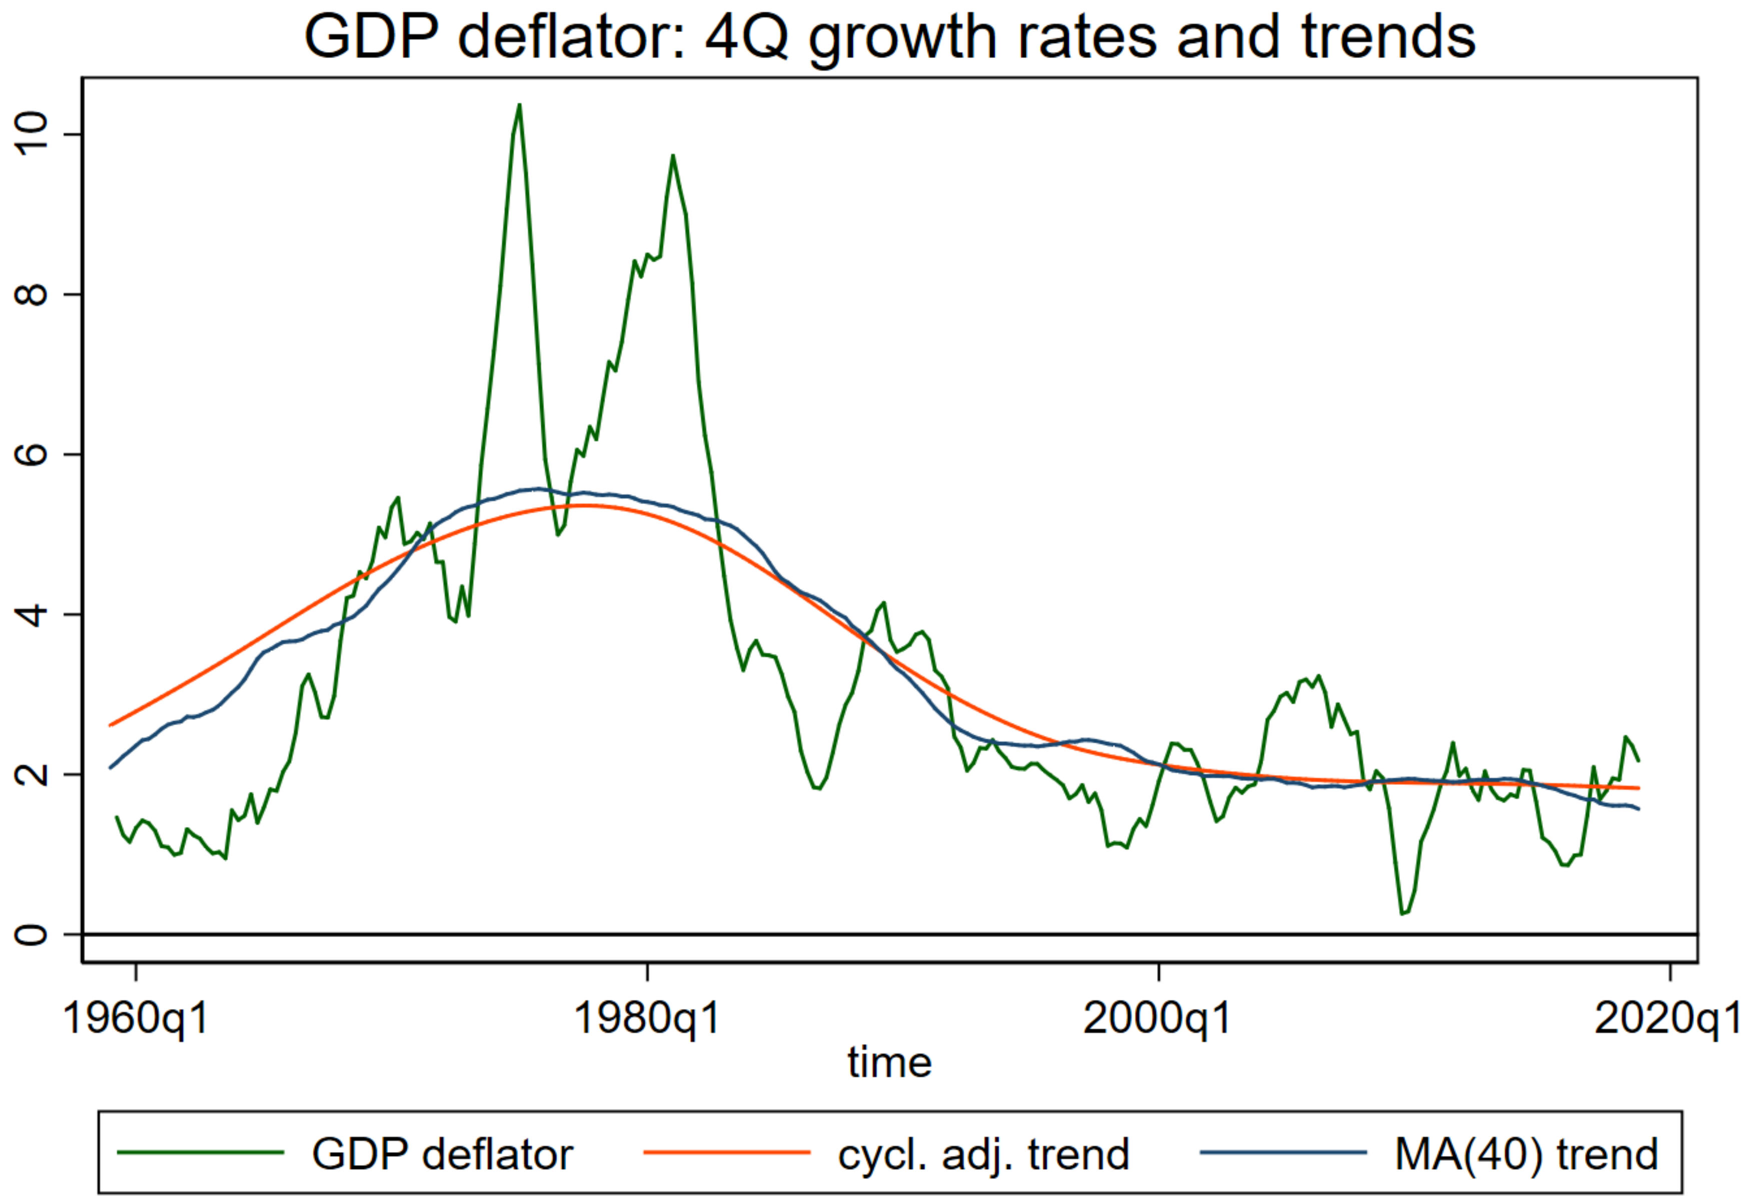

Sample mean, EWMA, and 75-year confidence intervals:

Statistic	%/yr
$\bar{x}_{1:T}$, 59-18	3.2
SE($\bar{x}_{1:T}$)	0.61
2019-2094	
EWMA, 2018:	2.2
67% CI	(1.4, 3.1)
67% MW (CPI)	(0.1, 5.3)
90% CI	(0.9, 3.6)

MW are Műller-Watson (2016) 67% confidence intervals using data from 1946-2014, 75-year projections

Comments & extensions

1. Get the variables right.
   a) Earnings are private production and supervisory workers, not SS earnings or total earnings
   b) Break down earnings further into (earnings/comp)×(comp/NGDP)

2. Adding series-specific knowledge or modeling. These are pure time series models – some series have structural elements that can be modeled.
   a) For example, the time series approach could be applied to the LFPR after subtracting out its aging component:

$$l_t = \Delta \ln(LFPR_t) = l_t^{aging} + \tilde{l}_t$$

   b) I have worked with per-capita GDP and earnings. Population growth could be modeled using these time series methods but other demographic or judgmental models might be preferred and would have less uncertainty.

Comments & extensions

3. Extending the formal prediction intervals to the local levels model?
   - Comparison of my CIs with Műller-Watson's suggests that making this formal extension probably won't increase the sampling uncertainty very much.
   - Also, the local levels model has pathologies when it is estimated – sometimes the "trend" simply replicates the series – so it cannot be applied mechanically.
     - This is why I imposed the EWMA smoother parameter, instead of estimating it.

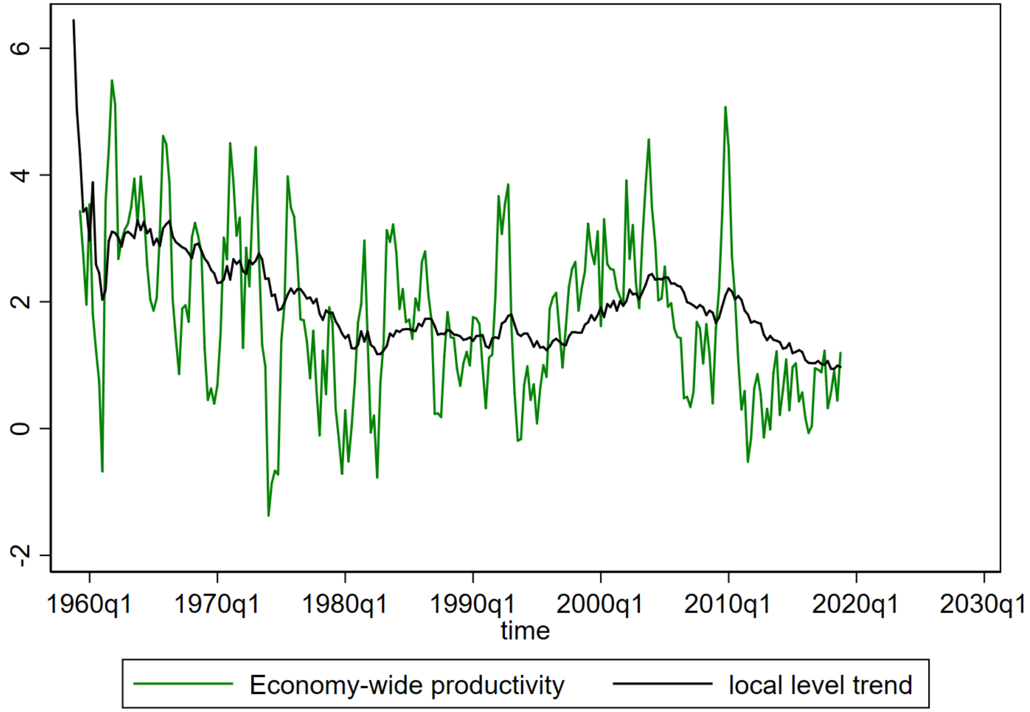

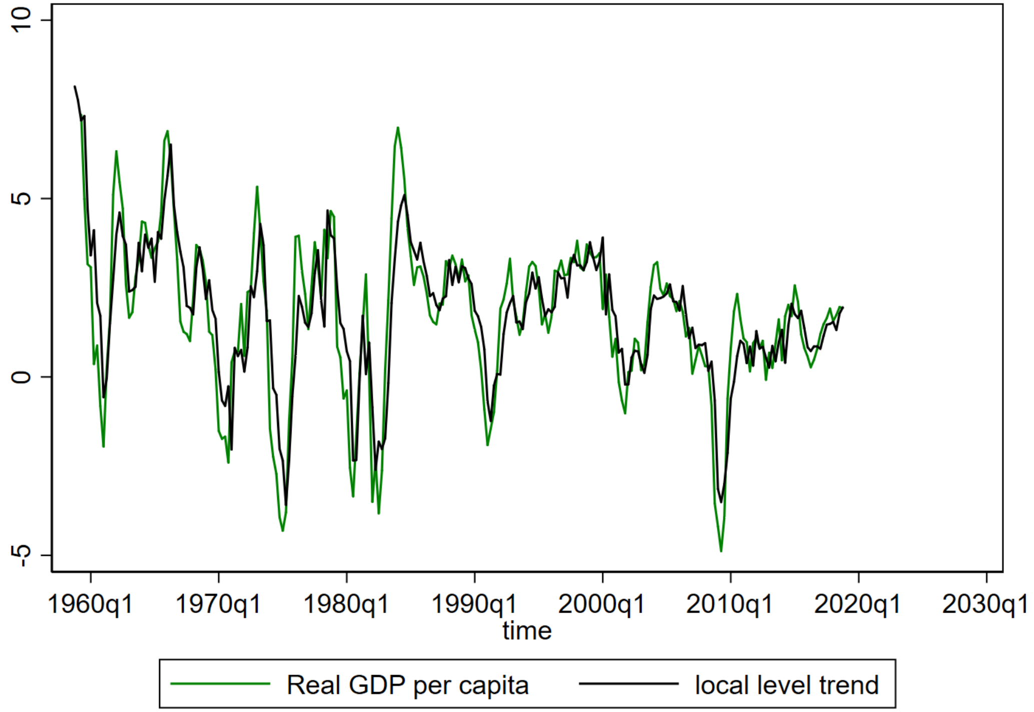

Comments & extensions

4. Multivariate distribution. If components are needed separately (for separate calculations) then they can be jointly drawn from the predictive distribution,

$$\bar{x}_{T+1:T+h} = \bar{x}_{1:T} + w, \text{ where } w \sim N\left[0, \left(\frac{1}{T} + \frac{1}{h}\right)\Omega\right]$$

where Ω is estimated by a multivariate HAC estimator.

- Drawing from the joint distribution addresses the problem of consistency across series in the "scenarios" approach.

5. Short-run dynamics. Formally the way this would be done is to estimate a Bayesian vector autoregression (VAR) with a diffuse prior; compute the predictive distribution; then sample from the predictive distribution. The predictive distribution captures both sources of uncertainty, forecast uncertainty and sampling uncertainty.
    - If some aspects of series are modeled separately (off-line), like LFPR aging component, then the VAR would be applied to the residuals.

Comments & extensions

6. Harder modeling (structural) extensions. For example:
   a) What is the effect of an increasing debt/GDP ratio on:
      - R*? (discussed below)
      - Capital accumulation?
   b) What are the interactions between demographics and key growth rates:
      - Labor productivity and aging (demographic composition of society).
      - Labor productivity and immigration
   c) Possibly use growth accounting identities instead of the current identity based on labor productivity (i.e. use capital/labor ratio, labor in efficiency units, TFP)
      - but none of these seem much easier to model structurally, or at least require substantial new modeling efforts (e.g. projections of education), and the literature to draw on is smaller.

Comments on the long-term real rate R*

Selected references: Laubach and Williams (2003), Del Negro et. al. (2017), Holsten, Laubach, and Williams (2017), Rachel and Summers (2019)

Rachel and Summers (2019, Fig. 1)

Figure 1: Real interest rates estimated from the inflation-linked bonds in advanced economies and in the United States

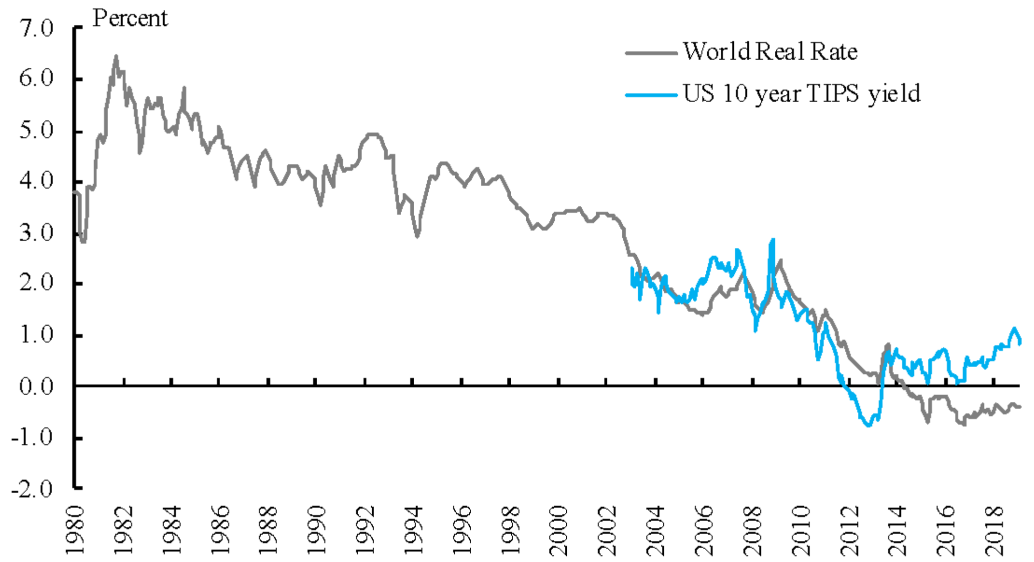

Note: The world real rate is calculated following the methodology in King and Low (2014): it is the average of interest rates on inflation-protected government debt securities across the G7 excluding Italy. Data are from DataStream and form an unbalanced panel. In particular, the Figure relies on the UK inflation-indexed gilts in the early part of the sample. The US TIPS yield is the yield on a constant maturity 10-year Treasury Inflation-Indexed Security, retrieved from FRED, Federal Reserve Bank of St. Louis (code DFII10).

Del Negro et. al. (2017, Fig. 1),

Figure 1. The Low-Frequency Component of r_t^* in the VAR and DSGE Models, 1960–2016[a]

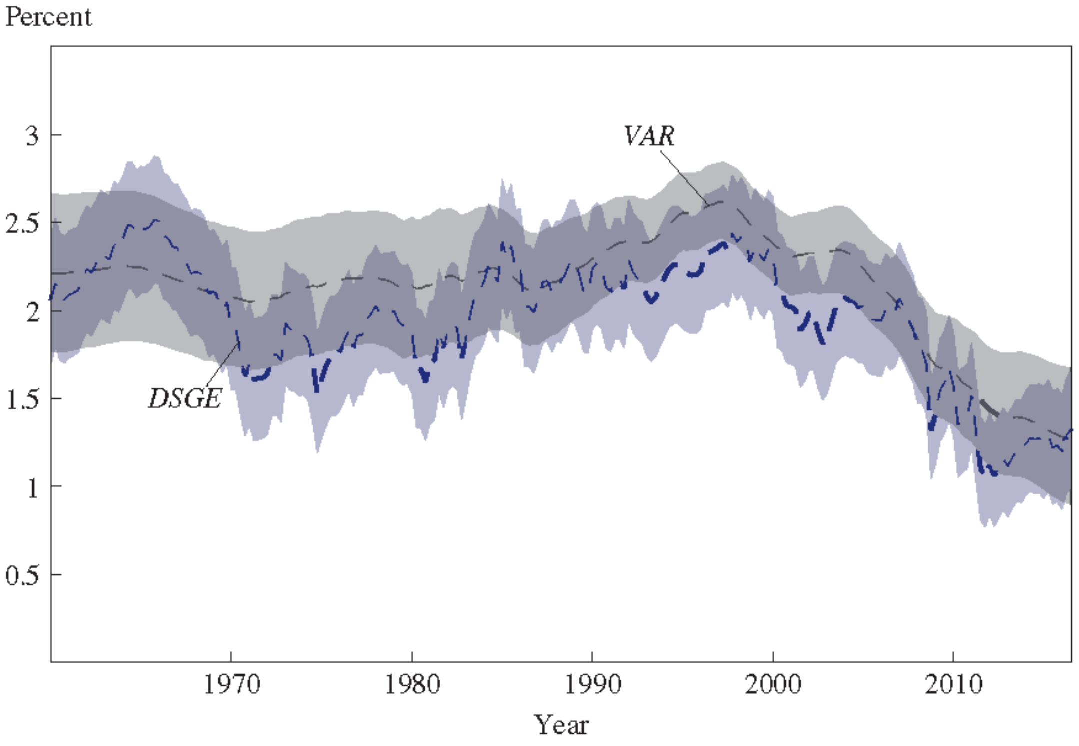

Source: Authors' calculations.

a. For each trend, the dashed line is the posterior median, and the shaded area shows the 68 percent posterior coverage interval for the estimate of the low-frequency component.

Summary

1. Uncertainty about long-term growth reflects (i) uncertainty about the future (forecast uncertainty) and (ii) uncertainty about population growth rates (estimation uncertainty)

2. This uncertainty can be captured rather simply just by focusing on long-term growth rates and sampling from a distribution using the long-run variance (i.e., using HAC standard errors)

3. Many of these series (in growth rates) contain long-term trends, and at least for some, it isn't at all clear that using a long historical average is the best jumping off point (e.g. LFPR, weekly hours).

4. A multivariate version of this provides internally consistent forecast uncertainty.

5. A natural reaction is that the variability of the postwar period was so large that we should not expect such variability in the future. For example, the entry of women into the labor force has already happened.
   - What about AI and robots? Gobalization and China? Climate change?